AL/CF-TR-1996-0100

ARMSTRONG LABORATORY

EFFECTS OF SAR IMAGE COMPRESSION ON IMAGE INTERPRETABILITY AND DETECTION PERFORMANCE

Judi E. See

LOGICON TECHNICAL SERVICES, INC.
P.O. BOX 317258
DAYTON, OH 45437-7258

Gilbert G. Kuperman

CREW SYSTEMS DIRECTORATE
HUMAN ENGINEERING DIVISION
WRIGHT-PATTERSON AFB OH 45433-7022

JUNE 1996

INTERIM REPORT FOR THE PERIOD OCTOBER 1995 TO JUNE 1996

AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OHIO 45433-6573

NOTICES

When US Government drawings, specifications, or other data are used for any purpose other than a definitely related Government procurement operation, the Government thereby incurs no responsibility nor any obligation whatsoever, and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication or otherwise, as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

Please do not request copies of this report from the Armstrong Laboratory. Additional copies may be purchased from:

National Technical Information Service
5285 Port Royal Road
Springfield, Virginia 22161

Federal Government agencies and their contractors registered with the Defense Technical Information Center should direct requests for copies of this report to:

Defense Technical Information Center
8725 John J. Kingman Road, Suite 0944
Ft. Belvoir, Virginia 22060-6218

TECHNICAL REVIEW AND APPROVAL

AL/CF-TR-1996-0100

This report has been reviewed by the Office of Public Affairs (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Julio B Cor

for **KENNETH R. BOFF,** Chief
Human Engineering Division
Armstrong Laboratory

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
June 1996

3. REPORT TYPE AND DATES COVERED
Interim Oct 1995 - Jun 1996

4. TITLE AND SUBTITLE
Effects of SAR Image Compression on Image Interpretability and Detection Performance

5. FUNDING NUMBERS
C: F41624-94-D-6000
PE: 62202F
PR: 7184
TA: 10
WU: 44

6. AUTHOR(S)
Judi E. See, PhD, Logicon Technical Service, Inc.
Gilbert G. Kuperman, Armstrong Laboratory

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Logicon Technical Services, Inc.
P.O. Box 317258
Dayton, OH 45437-7258

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Armstrong Laboratory, Crew Systems Directorate
Human Engineering Division
Human Systems Center
Air Force Materiel Command
Wright-Patterson AFB OH 45433-7022

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
AL/CF-TR-1996-0100

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution is unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The present study was designed to compare the relative impact of three image compression techniques on the interpretability of synthetic aperture radar (SAR) imagery: multiresolution encoding (MRES), Joint Photographic Experts Group (JPEG), and vector quantization (VQ). These three techniques were examined at three compression ratios (8:1, 16:1, and 32:1) and two image resolutions (fine and coarse) to determine which is most suitable for SAR imagery with respect to image interpretability and utility for target acquisition. Subjective ratings of image interpretability were obtained from six experienced image analysts for original and compressed SAR images via the Radar National Imagery Interpretability Rating Scale (RNIIRS). The subjective ratings were further supplemented with objective indicators of target detection performance (i.e., hits and false alarms) to permit a comprehensive assessment of each compression technique. For each dependent variable, different scores representing the difference in value for comparable uncompressed and compressed images were computed and analyzed. The results indicated that MRES was most similar to the uncompressed imagery in terms of both subjective and objective measures of effectiveness and may currently be a more suitable technique for SAR image compression than either JPEG or VQ.

14. SUBJECT TERMS
Image compression; Synthetic Aperture Radar (SAR); Image interpretability; Target detection

15. NUMBER OF PAGES
50

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UNLIMITED

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

This page intentionally left blank.

PREFACE

This effort was conducted by the Crew Systems Integration Branch, Human Engineering Division, of the Armstrong Laboratory (AL/CFHI), Wright-Patterson Air Force Base, Dayton, Ohio. The project was completed under Work Unit 71841044, "Crew-Centered Aiding for Advanced Reconnaissance, Surveillance, and Target Acquisition." Logicon Technical Services, Inc. (LTSI), Dayton, Ohio, provided support under contract F41624-94-D-6000, Delivery Order 0004. Mr. Donald Monk was the Contract Monitor.

The authors wish to acknowledge Mr. Raymond Withman of the Air Force Wright Laboratory for providing the imagery used in this effort. In addition, this work could not have been completed without the participation of the image analysts from the National Air Intelligence Center (NAIC) at Wright-Patterson Air Force Base in Dayton, OH.

The authors would also like to recognize the contributions of the following LTSI personnel throughout the duration of the project: Ms. Elisabeth Fitzhugh, who conducted the literature search; Mr. Brett Kottman, who installed the imagery and resolved subsequent display and processing problems; Ms. Jenny Huang, who developed the software to control image presentation and data collection in the experiment; Ms. Jacquelyn Grody, who helped test the software and procured photos and descriptions of the target vehicles; and Mrs. Iris Davis and Mr. Joseph Riegler, who provided helpful comments on earlier drafts of this report.

TABLE OF CONTENTS

TABLE OF CONTENTS, CONT'D.

LIST OF FIGURES

LIST OF TABLES

This page intentionally left blank.

INTRODUCTION

Image compression refers to the application of a set of techniques designed to enhance both storage efficiency and the speed of transmission of visual information by reducing the number of binary digits (bits) needed to represent an image. In general, the reduction is accomplished by coding the contents of an image and repackaging them in such a way as to achieve the desired bit rate. After transmission, the process is reversed to obtain a reconstructed representation of the original image. While image compression may improve storage and transmission efficiency, the conversion to a lower bit rate virtually always results in some loss in fidelity or an increase in image distortion. Consequently, it is necessary to identify and use those techniques that will enable the image to be reconstructed with as little degradation in operator performance during subsequent information extraction tasks as is possible for a particular type of imagery. Tolerable levels of degradation are task dependent; bomb damage assessment (BDA), for example, may be more tolerant of losses or distortions than would a target identification task.

One type of imagery that has seen recent widespread application of the techniques of image compression is the synthetic aperture radar (SAR) imagery that is used in numerous military applications because it may be acquired at long range and under adverse weather conditions. SAR imagery is exploited by human operators and by automatic target cuer/recognizer (ATC/ATR) systems during the detection and identification of military targets of interest. Because the imagery is collected by sensors on airborne platforms that have limited computational resources, it is typically transmitted to ground stations for processing. Effective use of the imagery requires that it be transmitted rapidly over the data link in a high quality format that can subsequently be readily interpreted during the target acquisition process. This need for rapid image transmission translates into a requirement for significant reductions in image bandwidth on the order of 10:1 or more via techniques that will simultaneously preserve the fidelity of the image. Hence, it is imperative to identify which techniques will adequately satisfy the requirements for SAR image compression.

Although a number of compression techniques are available, they have primarily been developed for use with optical (visible wavelength) imagery (e.g., still pictures and real-time television) and may not be entirely suitable for SAR imagery (Wharton, Gorman, Werness, & Gleason, 1993). First, SAR imagery does not possess the degree of statistical redundancy present in many other forms of imagery. Many current algorithms are successful because they are able to capitalize on the strong correlation between adjacent pixels that typically exists in optical images. Specifically, the information content of most optical images can be easily reduced without sacrificing image quality by eliminating or consolidating redundant and irrelevant pixel information. Because SAR imagery does not exhibit the high degree of pixel correlation that is representative of optical imagery, conventional compression algorithms must either be modified to take the random variations in SAR imagery into account or rejected altogether.

Second, unlike many other types of imagery where both low and high spatial frequency pixels may be sacrificed without compromising the overall utility and quality of the image, it is critical for SAR imagery that the high spatial frequency information contained in tactical targets be preserved for subsequent exploitation. Hence, the "equispaced rectangular grid" sampling strategy commonly used with optical imagery can seldom be applied to SAR imagery (Rabbani & Jones, 1991). With SAR imagery, it is necessary to sample more finely in target areas of high detail and assign more bits to those areas as opposed to coarser areas of low detail with fewer variations. Thus, any algorithms that may end up destroying high spatial frequency target information will not be suitable for the compression of SAR imagery.

SAR Compression Techniques

In an effort to select the algorithms that might be most suitable for use with SAR imagery, Wharton, Gorman, Werness, and Gleason (1993) applied a set of eight criteria to evaluate the effectiveness of 12 different classes of compression algorithms. These criteria included 1) achievable compression ratio, 2) achievable image quality, 3) suitability for SAR, 4) sensitivity to channel errors, 5) algorithm complexity, 6) algorithm maturity, 7) compatibility with standards, and 8) progressive transmission capability. The authors ranked the first three criteria as the most important in the evaluation of the efficacy of a compression algorithm. First, the capacity of an algorithm to achieve a high **compression ratio** is important for ensuring rapid transmission of imagery and availability of the data link. Second, unless the level of achievable **image quality** is high, the decompressed imagery will be virtually useless. Third, the algorithm must be equipped to handle the unique characteristics of **SAR imagery**.

The next two criteria, sensitivity to channel errors and algorithm complexity, were rated as having secondary and tertiary importance, respectively. **Sensitivity to channel errors** implies that the algorithm must incorporate some feature that will enhance its robustness to the bit errors (e.g., bit "flips," insertions, or deletions) that can occur during data transmission and seriously degrade image quality. **Algorithm complexity** refers to the type and number of computational requirements of the technique and the implications for image storage and the speed of transmission.

The final three criteria were judged to be equivalent in rank, but less crucial than previous criteria in the selection of a suitable algorithm. **Algorithm maturity** refers to established techniques whose strengths and limitations are well recognized as a consequence of widespread use. For this reason, conventional techniques may be preferred over more recently developed algorithms. However, as Wharton, Gorman, Werness, and Gleason (1993) indicate, powerful algorithms should not be rejected simply because they are not as mature as other approaches. In addition, maturity is typically defined in terms of experience with optical rather than SAR imagery. **Compatibility with standards** specified by the current National Imagery Transmission Format Standards (NITFS) for the transmission and storage of imagery would be a desirable feature of a compression technique since the imagery would be useable in a variety of domains (NITF, 1992). However, most of the standards have been defined in terms of the features of optical data and may therefore not be suitable for SAR imagery. Finally, **progressive transmission capability** received a low ranking because it is a nicety but not a necessity. This capability provides interactive viewing and transmission of an image; i.e, the receiver can browse through coarse representations of an image that have already been transmitted and choose areas that require subsequent transmission of finer resolution data.

Of the 12 classes of algorithms that were evaluated, seven of them were dropped from consideration because of failure to meet two or more of the three criteria judged to be most critical. Of the remaining compression techniques, the most suitable for SAR imagery were Joint Photographic Experts Group (JPEG; Pennebaker & Mitchell, 1993; Wallace, 1991), vector quantization (VQ; Gersho & Gray, 1992; Gray, 1984), and multiresolution (MRES) encoding (Mallat, 1989). All of these techniques are "lossy" as opposed to lossless methods for data compression. Lossless methods enable the original representation of the image to be fully reconstructed from the compressed data upon receipt. However, because none of the information in the image is compromised in lossless compression, only a limited amount of compression is possible. With lossy compression, on the other hand, some amount of information is lost or changed, and degradations in the reconstructed image are introduced relative to the original source image. Much higher compression ratios can be achieved, but at the expense of more distortion, which may or may not be operationally significant or visually apparent.

The general structure of a lossy compression system consists of three components: transformation, quantization, and encoding (Rabbani & Jones, 1991). First, decomposition or **transformation** of the data is completed to eliminate redundant information and provide a format that can be coded more efficiently. Second, **quantization** of the data is conducted to reduce the number of possible output symbols. It is this stage that introduces the loss in information but also makes higher compression ratios possible. The precise type and degree of quantization employed by a given technique affects both the transmission bit rate that can be achieved and the quality of the reconstructed image. Third, some form of symbol **encoding** process is used to represent the quantized data. The three techniques that Wharton, Gorman, Werness, and Gleason (1993) selected as the most appropriate for SAR imagery can be differentiated by the specific manner in which each of the three components is implemented. The following paragraphs provide brief general descriptions of each technique.

JPEG.

Of the three techniques, JPEG is the only one currently included in NITFS standards for the format and transmission of imagery. It is also currently used in several Department of Defense (DoD) systems (Withman, Cates, & Kotz, 1994). The technique was established under the auspices of the International Standardization Organization (ISO) by a group of experts working to develop a set of standards for image compression. Originally, the JPEG acronym stood for the group, but as the standard evolved and began to be applied, it came to be used to refer to the compression technique itself. The primary goal of the group was to develop a standard algorithm for the compression of continuous-tone, still-frame, monochrome and color images that would be applicable to areas as diverse as desktop publishing, graphic arts, color facsimile, photojournalism, and medical systems. As a result, JPEG contains a number of capabilities that make it suitable for a very wide range of applications. In fact, it has become the standard of comparison for judging new compression algorithms.

The JPEG standard consists of three main components: 1) a lossy baseline system that provides a simple and efficient algorithm that is sufficient for most applications; 2) a set of extended system features that allows the baseline system to meet a broader range of applications; and 3) an independent lossless method for applications that require reversible or bit-preserving compression.

In the JPEG baseline system, the image is first partitioned into blocks of 8x8 pixels, which are then compressed independently. **Transformation** of the data is accomplished via a discrete cosine transform (DCT), a decorrelation technique used to eliminate redundancies in the input. The DCT is a mature algorithm that is widely used and has the advantage of speed. However, it is not optimal for the low spatial correlations in SAR data and tends to produce blockiness, or visible discontinuities between adjacent blocks, in the image. Following transformation, a uniform **quantization** process is used to round each transformed coefficient to an integer value. The truncated DCT coefficients are then scaled by a factor derived from the coefficient amplitude that is just detectable by the human visual system. Because the performance of the human visual system is taken directly into account, this process makes the effects of the quantization less discernible during subsequent image exploitation (i.e., pixel values that are irrelevant to the human visual system can be discarded without causing any observable degradation in image quality). It is one of the critical factors behind the effectiveness of the JPEG algorithm. In fact, the largest compression gain for JPEG lies in its psychovisually-weighted quantization procedure (Lan, 1995). However, the quantization rules are established for a particular combination of viewing conditions and image content; if these conditions change, image distortion may become apparent.

Finally, **encoding** is achieved by means of differential pulse code modulation (DPCM), one of a general class of coding techniques known as predictive coding. In this scheme, information that has already been coded is used to predict future samples from the image so that the difference between the predicted and actual value can be encoded. Difference values, which tend to cluster around 0, especially when adjacent pixels are correlated, can be encoded much more efficiently with fewer bits than sample intensity values, which tend to exhibit much greater variability. Each difference value is assigned either a code word (Huffman coding) or a probability estimate (arithmetic coding), which is then converted to bits of compressed data. The JPEG encoding process serves to increase the compression ratio (up to 25:1 for SAR imagery) but is also very susceptible to channel errors. Further, it is not optimal for SAR since adjacent pixels tend not to be as highly correlated as in optical imagery.

A comparison of four different techniques conducted by Crocker and Cox (1995) supported the results of Wharton, Gorman, Werness, and Gleason's (1993) comprehensive survey by also demonstrating that JPEG could be used successfully with SAR imagery. Pixel by pixel comparisons of original and decompressed images indicated that the image quality of the reconstructed images was acceptable, as revealed by low mean square error (MSE) estimates and high signal to noise ratios (SNR) over six different compression ratios ranging from 2:1 to 64:1. This was true for both coarse resolution (wide-area search mode) and high resolution (spotlight mode) imagery. In addition, the reconstructed imagery continued to be exploitable by automatic detection devices, as revealed by an examination of the JPEG technique at ratios of 4:1, 8:1, and 16:1. Whereas JPEG remained comparable to the uncompressed baseline at a ratio of 4:1, however, performance degradations were observed at the higher ratios that are now routinely required in most applications.

<u>Vector quantization.</u>

The chief distinguishing feature of the VQ compression technique is that it operates on two-dimensional cells or vectors of data rather than on individual pixels. **Transformation** is achieved by decomposing the image into n-dimensional vectors to provide a format that can be coded more efficiently than scalars. During **quantization** and **encoding**, each vector is compared with a set of standard templates in a previously generated codebook, and the codeword identifying the best match is transmitted. The codeword providing the best match is the one that minimizes the distortion between the actual vector and the template, as assessed primarily by the MSE, a measure of the distance between the two vectors. The use of a codebook with relatively few codevectors compared to the number of possible image vectors is what enables compression to be obtained with VQ. Compression ratios of up to 40:1 are possible. At the receiver, the image is reconstructed using the templates in place of the original cells. The VQ technique generally permits the reconstruction of a good quality image that nevertheless has some blurring.

In essence, the principal goal of VQ is to represent the vectors of pixels in an image with a smaller subset of codebook vectors and still provide a good rendition of the data. The key step of the process is the selection of the optimum mapping from an input vector to one of a limited set of representative code vectors. The set of possible code vectors must be representative of the blocks in the actual image in order to minimize image distortion. A variety of VQ techniques, differentiated chiefly by the manner in which the codebook is selected and searched, are available. The version of VQ that was determined to be most suitable for SAR imagery is pruned tree-structured vector quantization (PTSVQ; Wharton, Gorman, Werness, & Gleason, 1993).

PTSVQ differs from other versions of VQ in terms of the codebook search strategy that is used. Codebooks are typically generated with a training set of images representative of the type of imagery to be encoded. In an attempt to find the codevector in the codebook that will yield the best match for the actual image vector, a full search of the resulting codebook can be performed. Because a full search can become computationally intensive, a tree-structured search, wherein only certain branches of the tree are examined for candidate codevectors, can be implemented instead. Distortion measurements are assigned to each encoding possibility, and the branch that produces the minimum distortion is ultimately selected. In PTSVQ, the codebook is derived from a large TSVQ that has been pruned back. Although this type of codebook allows for a more efficient search, it does make the PTSVQ technique more complex than either MRES or JPEG. In addition, PTSVQ utilizes variable rate quantization, in which the number of bits assigned to codewords is varied rather than constrained. That is, coding is done at a higher rate where needed (e.g., target areas of the image) and at a lower rate in portions of the image that are less crucial. Variable rate quantization is done to achieve higher ratios and better image quality, but it also means that PTSVQ is susceptible to channel errors.

Penrod and Kuperman (1993) demonstrated that VQ is more acceptable than two other commonly used techniques (block truncation coding and the discrete cosine transform) for SAR imagery. They examined both objective measures of image fidelity (MSE and SNR) and subjective ratings of image quality with a set of 23 SAR images. The analysis of the image fidelity measurements revealed that the VQ approach outperformed the other two compression techniques. Images subjected to VQ were more similar to their uncompressed counterparts in terms of both the MSE and the SNR. In addition, the subjective ratings provided by 12 image analysts indicated that VQ produced the least image distortion. However, these results apply only to the single compression ratio of 4:1 that was examined, a level that is quite low in comparison to the requirements of current applications.

Multiresolution.

In MRES coding, **transformation** is accomplished by decomposing the image into multiple uncorrelated resolution components. The image is first divided into a coarse version, followed by a fine scale version that contains only those additional details available at the finer resolution. This process is repeated until the image has been decomposed into successively finer resolution versions, each of which is uncorrelated with previous information. Coding of these composite frequency bands is more efficient than encoding the original signal, in part because the relative sensitivity of the human visual system to the various frequency bands can be taken into account when encoding the subsignals so that quantization noise will not be apparent in the reconstructed image. Thus, the more critical high frequency information contained in target areas of tactical imagery can be emphasized. The use of orthogonal frequency components also allows for progressive transmission of the image. **Quantization** and **encoding** are achieved through either scalar or vector quantization of the components.

Although MRES, JPEG, and VQ may all be more suitable for SAR imagery than other available techniques, additional investigations have further revealed that the MRES algorithms and JPEG may in fact be somewhat more effective than VQ. Withman, Cates, and Kotz (1994) examined the effects of image compression via the multiresolution wavelet approach, JPEG, and VQ during the "prescreening" stage of SAR ATC/ATR performance. During the prescreening phase, each pixel in the image is first processed to determine whether it is a potential target or clutter. Areas of the image designated as potential targets are then processed further in an attempt to reduce false alarm occurrences. The base set of uncompressed imagery used in the study consisted of 22 high resolution images (1024 x 1024 pixels). These images were compressed via each of the three techniques at ratios of 8:1, 16:1, and 32:1. Their results indicated that the first two techniques were somewhat more effective overall than VQ. The VQ approach tended to produce erratic performance effects in the ATC's hits and false alarms. As the compression ratio increased, performance deteriorated in some instances and inexplicably improved in others. Both JPEG and the multiresolution wavelet approach, on the other hand, exhibited more consistent performance degradation as the compression ratio increased. The results further indicated that performance was essentially unaffected by compression ratios of 8:1 or less. Even at this level, however, the multiresolution wavelet approach and JPEG were more effective than VQ, producing more hits and fewer false alarms.

Image Interpretability

While image compression offers many desirable features (e.g., reduced storage capacity, faster transmission, etc.), it may distort the image to the extent that its interpretability is adversely affected. Because the ultimate criterion for evaluating and accepting any algorithm is the opinion of human observers and their ability to use and interpret a reconstructed image, the subjective evaluation of image quality is critical for determining the effectiveness of a coding algorithm. Various types of rating scales tend to be used frequently to obtain such subjective evaluations of image interpretability.

The rating scale that is most widely used to assess the interpretability of optical imagery is the Image Interpretability Rating Scale (IIRS) adopted by the Air Standardization Coordinating Committee (ASCC, 1978) for the evaluation and reporting of photographic image quality (Itek, 1984). The IIRS provides 10 ratings that are based on the ground resolved distance (GRD; the sensor system resolution applied at the target) required to support various image analyst functions, including detection, counting, recognition, and identification of predefined targets of interest. The scale ranges from 0 (useless for image interpretation) through 9 (supporting the highest level of target identification). Rating Level 1 corresponds to a GRD greater than 9 m (30 ft), whereas Rating Level 9 corresponds to a GRD of less than 10 cm (4 in.).

The results of subjective assessments of various image compression techniques via the IIRS have further supported the suitability for SAR imagery of at least one of the three approaches identified as most appropriate by the objective evaluations described earlier. Vaughn (1987) used the IIRS to assess the effectiveness of VQ for reconnaissance imagery in comparison to two common scalar quantizers (DPCM and DCT). The images included in the evaluation were 50 digital infrared and black and white images. Two photo interpreters rated the interpretability of each uncompressed and compressed image with the IIRS. For the original images, average IIRS ratings ranged from 3.3 to 8.5. At the 4:1 compression ratio, the ratings for the three compression techniques were comparable to one another and lagged behind the ratings of the original imagery by only two-tenths of a point. However, at the 8:1 ratio, VQ was more effective than its competitors, exhibiting better preservation of ground resolution than the other schemes. Ratings for VQ compressed images at this ratio were only .1 lower than those for the original imagery, whereas ratings for the other two techniques were lower by nearly a full category (.8 lower on average).

Following the successful application of the IIRS, a structurally similar subjective interpretability rating scale was developed for use with SAR imagery, the Radar National Imagery Interpretability Rating Scale (RNIIRS). The RNIIRS is presented in the Appendix. Results similar to those found by Vaughn (1979) with the IIRS were obtained by Penrod and Kuperman (1993) using the RNIIRS. Twelve image analysts provided RNIIRS ratings for 23 SAR images that had been subjected to one of three compression techniques at a ratio of 4:1. The analysts' subjective ratings indicated that VQ produced less image distortion (mean RNIIRS rating = 2.71) than either the DCT ($\underline{M}$ = 2.69) or another commonly used technique known as block truncation coding ($\underline{M}$ = 2.52). Thus, the VQ compressed images were most comparable to the original imagery ($\underline{M}$ = 2.83).

Although no studies in which the RNIIRS has been used in conjunction with MRES and JPEG encoding could be located, one other relevant investigation used the IIRS to study the effects of compression ratio (4:1 and 2:1) on image interpretability for a single technique (Wilson & Kuperman, 1989). Thirty image analysts provided IIRS ratings for 42 tactical reconnaissance images (14 each of uncompressed, 4:1, and 2:1 imagery). The results indicated that the IIRS ratings did not differ as a function of compression ratio. The mean ratings for 4:1 and 2:1 compressed images were 4.73 and 4.70, respectively. However, overall ratings for compressed imagery were significantly lower than those for the original imagery by approximately .5 IIRS category units. As noted by Wilson and Kuperman (1989), a difference of this magnitude can reduce the level of exploitation (e.g., from identification to recognition) or the analyst's decision confidence (e.g., from "probable" to "possible"). In sum, their study and others have shown the IIRS and the RNIIRS to be useful tools for examining the effects of both the type and amount of image compression on image interpretability.

The Present Investigation

The purpose of the present study was to augment current information regarding the effects of image compression on the utility of SAR imagery. Wharton, Gorman, Werness, and Gleason's (1993) evaluation of various alternative compression techniques identified MRES, JPEG, and VQ to be the most suitable algorithms for SAR imagery. The eight criteria used in their assessment represent the qualities that make a given technique desirable to use with SAR imagery. The ultimate criterion, above and beyond these desirable qualities, however, is whether the compressed imagery can be exploited to the same extent as uncompressed imagery during target acquisition. The relative impact of each of the three compression techniques on image interpretability and utility has not yet been thoroughly explored. Consequently, the primary purpose of the present study was to fill this void.

Three compression techniques (MRES, JPEG, and VQ) were investigated at each of three compression ratios (8:1, 16:1, and 32:1) for both fine and coarse resolution SAR imagery. The chief goal of the current investigation was to determine which technique is most suitable for SAR imagery in terms of image interpretability across various compression ratios and resolutions within a single study. Subjective ratings of image interpretability were obtained from experienced image analysts for original and compressed SAR images via the RNIIRS. An additional goal was to supplement the subjective ratings with objective indicators of target detection performance (i.e., hits and false alarms) to permit a more comprehensive assessment of each compression technique. Whereas the effects of image compression on the detection performance of ATC/ATR devices has been examined previously, the impact on human performance has not. Accordingly, the image analysts in the current study were also asked to designate the location of each target in the original and compressed imagery in order that hits and false alarms could be derived and analyzed.

METHOD

Participants

The participants were six image analysts (one female) from the National Air Intelligence Center (NAIC) at Wright-Patterson Air Force Base in Dayton, OH. They had worked as image analysts for 1 to 9 years ($\underline{M} = 6.7$, $\underline{SD} = 2.7$). Five analysts used SAR imagery with some degree of frequency in their current assignment, and all of them had previously used image interpretability rating scales, including the RNIIRS. The participants ranged in age from 25 to 45 years ($\underline{M} = 33.8$, $\underline{SD} = 6.4$). All had normal or corrected-to-normal 20/20 visual acuity or better, as measured with a Snellen visual acuity chart.

Experimental Design

Two levels of image resolution (fine and coarse) were combined factorially with three image compression techniques (MRES, JPEG, and VQ) and three compression ratios (8:1, 16:1, and 32:1) to provide 18 conditions. Two additional conditions were provided by the presentation of uncompressed imagery at both the fine and coarse resolutions. There were 22 unique fine spot scenes and 19 unique coarse spot scenes. Each was presented 10 times (one uncompressed image plus the nine images provided by the 3 x 3 combination of techniques and ratios), for a total of 410 images.

Apparatus and Stimuli

The images used in the present study were a subset of the SAR imagery collected during the Chicken Little exercise at Camp Grayling, Michigan, in October 1992. Only those images that had been assigned "Good" quality ratings by experienced SAR image analysts were selected for subsequent compression; of those, 22 fine spot and 19 coarse spot images were included in the current study. The targets were 13 Soviet ground order battle (GOB) vehicles deployed in patterns designed to conform as closely as possible to Soviet military doctrine. The target site was an open area surrounded by a heavily forested background. Many of the targets were located completely in the open with no obscuration. In some cases, however, obscuration of individual targets ranged from moderate to heavy.

The targets consisted of six types of armored personnel carriers (APCs--BMP-1, BMP-2, BM-21, BTR-60, BTR-70, and BTR-80), two tanks (T-72 and T-62), two trucks (TZM and Zil-131), two self-propelled howitzers (2S1 and 2S3), and an anti-aircraft gun (ZSU-23/4). Complete descriptions, line drawings, and specifications for each target can be found in the ground truth report (Lewis, 1993). Each fine spot image contained between 1 and 11 targets, providing a total of 133 targets in the set of 22 images. Each coarse spot image contained a minimum of 10 and a maximum of 27 targets, for a total of 395 targets in the set of 19 images. The distribution of targets depicted in Table 1 shows that the T-72 tank was the most prevalent target type in both the fine and coarse spot imagery.

Each unique image was compressed via the MRES, JPEG, and VQ techniques at ratios of 8:1, 16:1, and 32:1. Details of the specific algorithms that were used can be found in Wharton, Gorman, Werness, and Gleason (1993). The MRES algorithm that was used was referred to as multiresolution/subband coding with hybrid scalar/vector quantization, and the VQ algorithm was a PTSVQ. All compressed and uncompressed images were 1024 x 1024 pixels in size. The imagery was presented on a high-resolution Silicon Graphics computer monitor whose screen size was 15 1/8 in. wide by 11 1/2 in. tall. A supplementary keypad was attached for use in the experiment. The keypad contained 10 keys labeled 0 through 9 for the entry of image interpretability ratings provided by the Radar National Imagery Interpretability Rating Scale (RNIIRS). Although some applications of this rating scale have permitted non-integer ratings (e.g., .5, 1.5, etc.), only integers were used in this study. Two keys labeled "Ready" and "Enter" were also used during the experiment. In addition, the first (left) and third (right) buttons of the computer mouse were used during the designation and undesignation, respectively, of potential target objects in each image.

Table 1

The Distribution of Targets in Fine and Coarse Spot Imagery

	Percent Occurrence	
Target Type	FINE SPOT	COARSE SPOT
APCs		
BMP-1	3	3
BMP-2	10	12
BM-21	9	6
BTR-60	9	7
BTR-70	3	2
BTR-80	2	3
Tanks		
T-72	29	26
T-62	3	10
Trucks		
TZM	4	3
Zil-131	6	6
Self-propelled howitzers		
2S1	14	12
2S3	1	3
Anti-aircraft gun		
ZSU-23/4	7	7
Total	100% (N=133)	100% (N=395)

Procedure

The experiment was conducted in the Crew-Aiding and Information Warfare Analysis Laboratory (CIWAL) at Armstrong Labs, Wright-Patterson Air Force Base, OH. Each individual participated in two separate sessions on two different days. All of the coarse spot imagery was presented during one three-hour session (190 images), and all of the fine spot imagery was presented during a two-hour session (220 images). The order of occurrence of the two sessions was randomly determined for each participant. Three individuals completed the coarse spot session first, and three completed the fine spot session first. Upon arriving for the experiment, participants were tested for normal visual acuity with the Snellen visual acuity chart in typical indoor lighting conditions. They were then escorted to the room within CIWAL in which the experiment was conducted and were seated before the computer used during completion of the experimental task. The computer was placed on a desk with the monitor at approximately eye level. The standard keyboard was replaced with the supplementary keypad and mouse arrangement. During each session, the only light in the room was provided by a desk lamp, which was positioned so as to provide ample lighting but minimize glare on the computer monitor.

Participants were first provided with a general description of the experimental task before being asked to read and sign a consent form, which provided additional details. Following their consent to participate, individuals completed a brief questionnaire, which primarily assessed their familiarity with SAR imagery and with image interpretability rating scales. The purpose of the current study was described next as were the details regarding the imagery, the targets to be detected, and the RNIIRS. Individuals were told that thirteen types of target vehicles were present in the imagery and that obscuration could range from none to heavy. Black-and-white photographs of the targets (from the Janes reference books) and brief descriptions of each, including their dimensions, were provided. To avoid prolonging the session, participants had been asked to review the RNIIRS prior to their arrival. Any questions regarding its usage in the current experiment were answered before the session began. A copy of the RNIIRS was available for the duration of the experiment. Individuals were told that they would be asked to evaluate 190 coarse spot images during a three-hour session and 220 fine spot images during a separate two-hour session. They were informed that they would be asked to rate the interpretability of each image using the RNIIRS and that they would be asked to detect and designate all targets in each coarse spot image (containing 10 - 27 targets) and each fine spot image (containing 1 - 11 targets).

The procedure to be followed on each trial was described next. Each trial began with a message in the center of the computer screen instructing the participant to press the "Ready" key on the keypad placed in front of the computer. An image appeared on the monitor, and a brief message instructed the individual to rate the image's interpretability first by pressing a number from 0 through 9 followed by the "Enter" key. The rating could be changed as many times as necessary prior to pressing "Enter." Only the last number pressed before entry was accepted and recorded as the image interpretability rating. After the rating had been entered, the image remained on the screen and a new message appeared instructing the participant to designate all targets in the image. The participant was instructed to center the cursor over a target and click with the first mouse button to designate. Each designated object was overlaid with a small white circle, numbered according to the order of occurrence. The third mouse button could be used to undesignate any previous target designation. An undesignation resulted in removal of the overlay and renumbering of the remaining symbology to reflect the change. The individual was instructed to press "Enter" when done to submit those designations remaining on the screen. For each image, the numbers of correct (hits) and incorrect (false alarms) designations were recorded. Following entry of the target designations, the "Ready" prompt again appeared in the center of the screen. A flow chart of the sequence of events during each trial is depicted in Figure 1.

Because the time allotted for ratings and target designations was unlimited, the task was completely self-paced. Further, participants were not only instructed at specific times during each session to take breaks but were also encouraged to take additional breaks as needed. Messages instructing the participant to take a 10-minute break appeared on the computer at the conclusion of the first third and the second third of the coarse spot session and at the end of the first half of the shorter fine spot session. Individuals were informed that they could, however, take a break at any time during the session at their own discretion or no break at all if they chose. The duration of the average coarse spot session was 3 hours (SD = 30 min.), and the duration of the average fine spot session was 2 hours (SD = 20 min.).

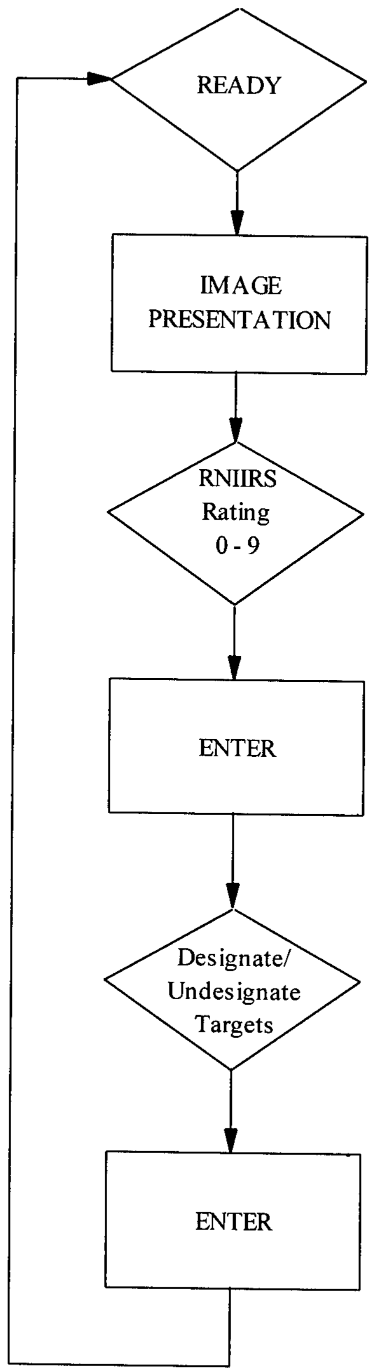

Figure 1. The sequence of events during each experimental trial.

RESULTS

The principal goal of the present study was to examine the impact of image compression on observers' ability to utilize SAR imagery for target acquisition by analyzing both subjective and objective measures. Observers' RNIIRS ratings provided the subjective measure, and their correct and incorrect designations provided an objective indicator of the effects of compression on image utility. While it is useful to examine the raw scores for each compression technique, the most informative indicator of the effectiveness of a given technique is the comparison between its scores and those for the uncompressed imagery. Consequently, the primary dependent variables of interest in the current investigation were *difference* scores, computed as the difference between uncompressed and compressed scores. In this type of analysis, the uncompressed score serves as a baseline against which all other scores are compared. In the sections that follow, the raw untransformed scores in each experimental condition are described first to provide the absolute magnitudes of RNIIRS ratings, hits, and false alarms in each condition. Next, procedures for computing the transformed scores are described. Finally, the difference scores are tabulated and the outcomes of the statistical analyses of these data are covered.

Raw Score Data

RNIIRS ratings.

Mean raw score RNIIRS ratings by resolution, compression technique, and compression ratio are depicted in Table 2. As can be seen in the table, fine resolution images received higher RNIIRS ratings than coarse resolution images, verifying that higher resolutions do indeed support the detection of finer details. Further, the image interpretability ratings were lower for all three compression techniques as compared to those for the original (uncompressed) imagery. Of the three techniques, MRES had the highest ratings and was therefore most comparable to the uncompressed imagery. Ratings for JPEG were lowest. Finally, the means in Table 2 indicate that ratings tended to decrease as the compression ratio increased. The overall RNIIRS rating of 3.22 signifies that the image analysts judged the quality of the imagery used in the present study to be useful primarily for target detection functions but not for classification or identification.

Table 2

Means and Standard Deviations (in Parentheses) for RNIIRS Ratings

	COARSE RESOLUTION				FINE RESOLUTION				
	8:1	16:1	32:1	Mean	8:1	16:1	32:1	Mean	MEAN
JPEG	3.06 (0.48)	2.67 (0.28)	1.97 (0.50)	2.57 (0.36)	3.76 (1.18)	3.67 (1.07)	3.10 (1.16)	3.51 (1.11)	3.07 (0.61)
MRES	3.10 (0.52)	2.94 (0.47)	2.67 (0.62)	2.90 (0.52)	3.76 (1.07)	3.73 (1.15)	3.61 (1.21)	3.70 (1.14)	3.33 (0.67)
VQ	2.72 (0.42)	2.73 (0.42)	2.43 (0.54)	2.62 (0.44)	3.61 (1.18)	3.60 (1.22)	3.52 (1.17)	3.58 (1.19)	3.14 (0.66)
ORIGINAL	3.25 (0.51)				3.84 (1.21)				3.57 (0.76)
Mean	2.75 (0.42)				3.62 (1.14)				3.22 (0.64)

Hits.

In order for the observer's designation to be classified as a hit, it had to be located no more than one "target" length away from the vehicle; i.e., the distance of the designation from the center of the target could be no greater than the approximate size of the largest target at each resolution. The percentage of hits was computed as a ratio of the number of correct detections to the total number of targets. Mean percentages of hits by resolution, compression technique, and compression ratio are portrayed in Table 3. As can be seen in Table 3, the percentage of correctly detected targets was higher at the fine resolution than at the coarse resolution. Although the percentage of hits was consistently lower for each compression technique as compared to the original imagery, the discrepancy was fairly small in each case. Finally, the percentage of hits declined as the compression ratio increased. The mean percentage of correct detections for the set of imagery was 75%. Thus, even though the imagery was judged to be sufficient for target detection, target detection performance itself was rather poor. Altogether, 25% of the targets present in the imagery were missed.

Table 3

Means and Standard Deviations (in Parentheses) for Percentages of Hits

	COARSE RESOLUTION				FINE RESOLUTION				
	8:1	16:1	32:1	Mean	8:1	16:1	32:1	Mean	MEAN
JPEG	66 (5.2)	65 (5.1)	60 (4.6)	64 (4.7)	87 (4.8)	86 (3.6)	84 (3.7)	86 (3.9)	75 (3.3)
MRES	67 (4.4)	66 (3.6)	62 (5.8)	65 (4.3)	86 (4.6)	86 (4.0)	83 (4.4)	85 (3.8)	76 (2.9)
VQ	65 (3.8)	61 (5.6)	60 (5.7)	62 (5.0)	83 (4.4)	83 (4.0)	83 (4.6)	83 (4.0)	73 (3.1)
ORIGINAL		68 (5.2)				86 (4.4)			78 (3.8)
Mean	64 (4.6)				85 (3.9)				75 (3.1)

False alarms.

False alarms consisted of designations that lay outside the specified range for correct detections; i.e., designations of bright returns in the scene that were not target objects. Mean numbers of false alarm responses by resolution, compression technique, and compression ratio are presented in Table 4. Unlike the hit responses, false alarm responses could not be converted to a percentage score since the number of false alarm possibilities in a given SAR image is potentially limitless. That is, whereas targets are distinct objects that are easily quantified, nontargets represent *any* aspect of the scene that might possibly be mistaken for a target and may not even be tangible "objects" at all. The means in Table 4 indicate that the average number of false alarms for the set of imagery was only one. This outcome is not surprising since the image analysts were informed of the minimum and maximum numbers of targets in the imagery prior to each session. Fewer false alarms were committed at the fine resolution than at the coarse resolution. In addition, the quantity of false alarms was similar for compressed and uncompressed imagery. Finally, the tendency to commit false alarms was only minimally affected by compression ratio.

Table 4

Means and Standard Deviations (in Parentheses) for Numbers of False Alarms

	COARSE RESOLUTION				FINE RESOLUTION				
	8:1	16:1	32:1	Mean	8:1	16:1	32:1	Mean	MEAN
JPEG	1.6 (0.4)	1.7 (0.4)	1.4 (0.3)	1.6 (0.2)	0.4 (0.3)	0.4 (0.3)	0.3 (0.3)	0.4 (0.3)	0.9 (0.2)
MRES	1.7 (0.5)	2.3 (0.8)	2.0 (0.6)	2.0 (0.6)	0.4 (0.3)	0.4 (0.3)	0.4 (0.4)	0.4 (0.3)	1.2 (0.4)
VQ	1.9 (0.3)	2.0 (0.5)	2.0 (0.8)	2.0 (0.4)	0.4 (0.3)	0.4 (0.3)	0.3 (0.3)	0.4 (0.3)	1.1 (0.3)
ORIGINAL	2.3 (0.8)				0.3 (0.3)				1.2 (0.4)
Mean	1.9 (0.4)				0.4 (0.3)				1.1 (0.3)

Difference Scores

The raw scores from which the means in Tables 2 through 4 were computed were used to derive difference scores for the RNIIRS ratings, hits, and false alarms. In each case, the uncompressed score served as the baseline from which the comparable compressed score was subtracted for each unique image. For each dependent variable, this calculation provided a total of 22 difference scores per subject at the fine resolution and 19 at the coarse resolution for each of the nine compression techniques (MRES, JPEG, and VQ at the 8:1, 16:1, and 32:1 ratios). The mean difference score for each subject at each technique was then computed by averaging the scores for the 22 images at the fine resolution and those for the 19 images at the coarse resolution. Within each resolution, this calculation resulted in 9 mean difference scores per subject, one for each technique. The overall means for each technique were then derived by averaging across the scores for the 6 subjects. All difference scores were analyzed by means of separate 2(resolution) x 3(compression technique) x 3(compression ratio) repeated measures analyses of variance. The significance of F ratios involving compression technique and compression ratio was determined by the Huynh-Feldt epsilon adjustment procedure (Huynh & Feldt, 1970; 1976). Follow-up tests consisted of post hoc t-tests, with the level of significance adjusted by means of the Bonferroni procedure to control the Type I error rate. When only one post hoc test was conducted, the level of significance was set at the usual alpha of .05. For each set of multiple post hoc tests, the overall alpha was set at .20. Thus, in those cases where follow-up testing entailed three comparisons, the resulting level of significance for each individual test was .07. For nine comparisons, the level of significance for each individual test was .022. In essence, this procedure requires a more stringent criterion for significance as the number of comparisons increases.

RNIIRS ratings.

Mean RNIIRS rating difference scores by resolution, compression technique, and compression ratio appear in Table 5. For these difference scores, a value of 0 indicates that the uncompressed and compressed images received similar ratings. Positive values signify that the uncompressed imagery received higher image interpretability ratings, whereas negative values indicate that the compressed images were judged to be of higher quality. As can be seen in the table, across all resolutions, techniques, and ratios, the uncompressed imagery received higher RNIIRS ratings than the compressed imagery. The gap between uncompressed and compressed imagery was greater at the coarse resolution than at the fine resolution, signifying that compression had the greatest impact on the interpretability of coarse resolution imagery. The means in the table also indicate that the ratings for MRES were most similar to those for the original imagery, whereas those for JPEG were least similar. Further, the gap between uncompressed and compressed imagery tended to widen as the compression ratio increased.

Table 5

Means and Standard Deviations (in Parentheses) for RNIIRS Rating Difference Scores

	COARSE RESOLUTION				FINE RESOLUTION				
	8:1	16:1	32:1	Mean	8:1	16:1	32:1	Mean	MEAN
JPEG	0.19 (0.25)	0.59 (0.34)	1.28 (0.52)	0.69 (0.31)	0.08 (0.15)	0.17 (0.24)	0.74 (0.46)	0.33 (0.22)	0.50 (0.24)
MRES	0.15 (0.20)	0.32 (0.29)	0.59 (0.38)	0.35 (0.26)	0.08 (0.20)	0.11 (0.16)	0.23 (0.12)	0.14 (0.11)	0.24 (0.16)
VQ	0.54 (0.30)	0.53 (0.30)	0.82 (0.51)	0.63 (0.36)	0.23 (0.15)	0.24 (0.11)	0.32 (0.24)	0.26 (0.15)	0.43 (0.21)
Mean	0.29 (0.22)	0.48 (0.27)	0.90 (0.33)		0.13 (0.12)	0.17 (0.11)	0.43 (0.16)		
Mean	0.56 (0.26)				0.24 (0.10)				0.39 (0.15)

The analysis of variance of the scores in Table 5 revealed that the main effects for resolution, $F(1,5) = 9.85$, $p < .03$, and ratio, $F(2,10) = 40.34$, $p < .0002$, were statistically significant. Post hoc correlated t-tests revealed that all three ratios were significantly different from one another. Although the main effect for compression technique was not significant ($p > .05$), its interactions with resolution and ratio were: $F(2,10) = 4.37$, $p < .043$ for the Compression Technique x Resolution interaction and $F(4,20) = 6.73$, $p < .02$ for the Compression Technique x Ratio interaction. The Resolution x Ratio interaction was also statistically significant, $F(2,10) = 7.60$, $p < .02$. However, the three-way interaction did not attain statistical significance, $F(4,20) = .86$, $p > .05$.

The nature of the interaction between compression technique and resolution is portrayed graphically in Figure 2. The main plot in Figure 2 shows the mean RNIIRS difference scores for each technique at the coarse and fine resolutions, and the smaller inset shows the resolution effect for MRES, VQ, and JPEG (i.e., the difference between the scores at the coarse and fine resolutions for each technique). As portrayed by the two plots in the figure, the RNIIRS difference scores were consistently lower at the fine resolution than at the coarse resolution for all three compression techniques; however, the magnitude of the disparity between the two resolutions was smaller for MRES than for either VQ or JPEG. That is, of the three techniques, MRES had the smallest impact on the image interpretability rating at both the coarse and fine resolutions. A post hoc t-test revealed that the resolution effect for MRES was significantly smaller than that for VQ and JPEG, $t(10) = 2.53$, $p < .05$.

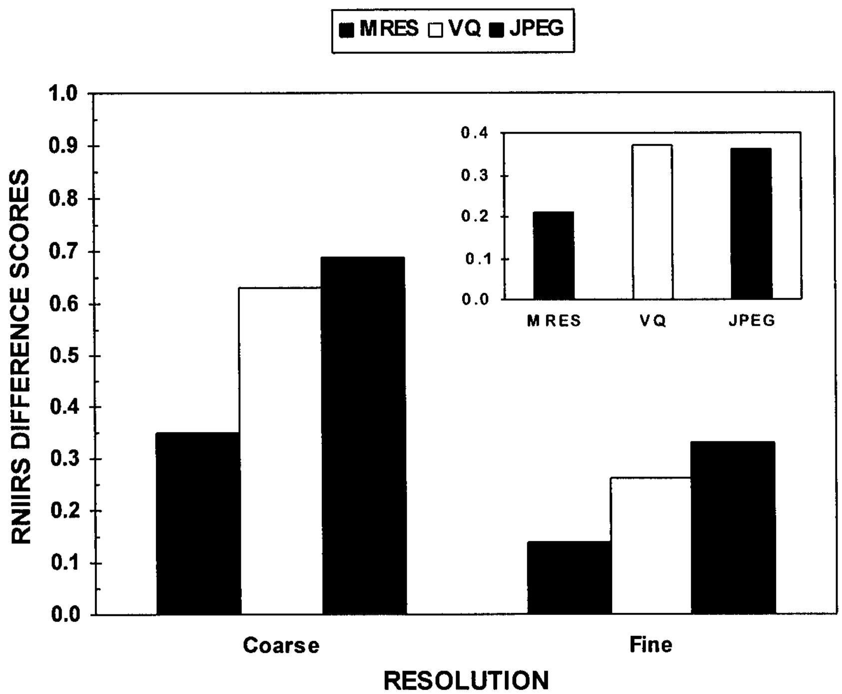

Figure 2. Mean RNIIRS difference scores for each compression technique at the coarse and fine resolutions. (Note: The inset shows the difference between the resolutions for each technique.)

The Resolution x Ratio interaction is depicted in Figure 3. The main plot in the figure portrays the mean RNIIRS difference score for ratios of 8:1, 16:1, and 32:1 at the coarse and fine resolutions. The inset depicts the resolution effect at each compression ratio (i.e., the difference between the scores at the coarse and fine resolutions for each ratio). As can be seen in the figure, the RNIIRS difference scores were consistently larger at the coarse resolution than at the fine resolution for all three compression ratios. However, the difference between the two resolutions was minimal at the 8:1 ratio but sizable at the higher compression ratios, an effect that was statistically significant, $t(10) = 2.75$, $p < .05$.

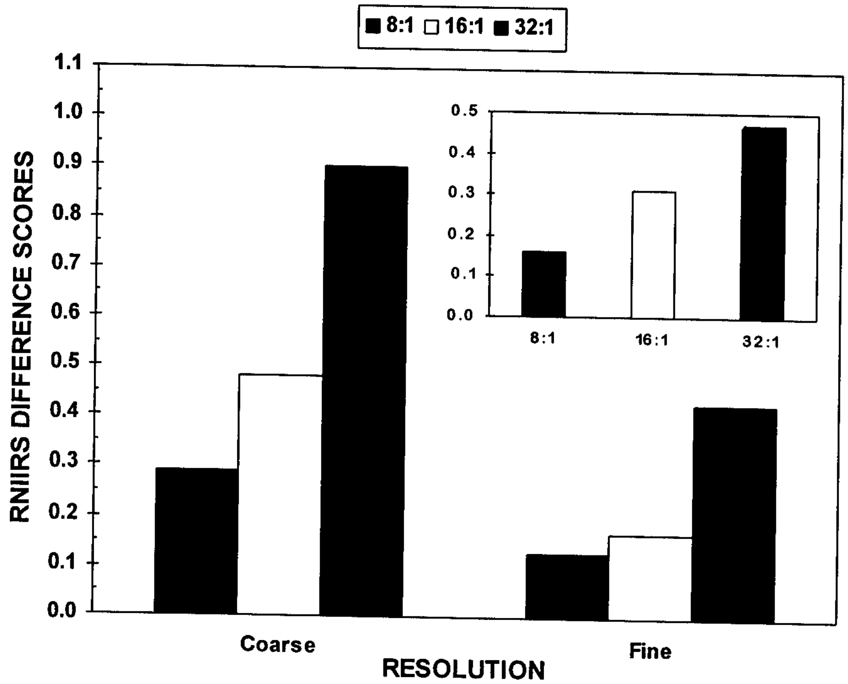

Figure 3. Mean RNIIRS difference scores for each compression ratio at the coarse and fine resolutions. (Note: The inset shows the difference between the resolutions at each compression ratio.)

The Compression Technique x Ratio interaction is portrayed in Figure 4. For MRES, the RNIIRS difference score increased progressively as the compression ratio increased, indicating that the gap between the ratings for uncompressed and compressed imagery increased with the compression ratio for MRES. For VQ, the difference score remained stable until the ratio reached 32:1. For JPEG, the difference score increased progressively as ratio increased. The sharp increase from the ratio of 16:1 to 32:1 indicates that imagery subjected to a JPEG compression will be interpretable at lower ratios but not at the higher compression ratios that are frequently required for image storage and transmission. Post hoc correlated t-tests were conducted to determine whether there were significant differences among the three ratios within each compression technique. For MRES, the 8:1 ratio differed from 16:1 and from 32:1; however, the latter two ratios were not significantly different. For VQ, none of the ratios differed from one another. Finally, for JPEG, the ratios of 8:1 and 16:1 differed from 32:1 but not from each other.

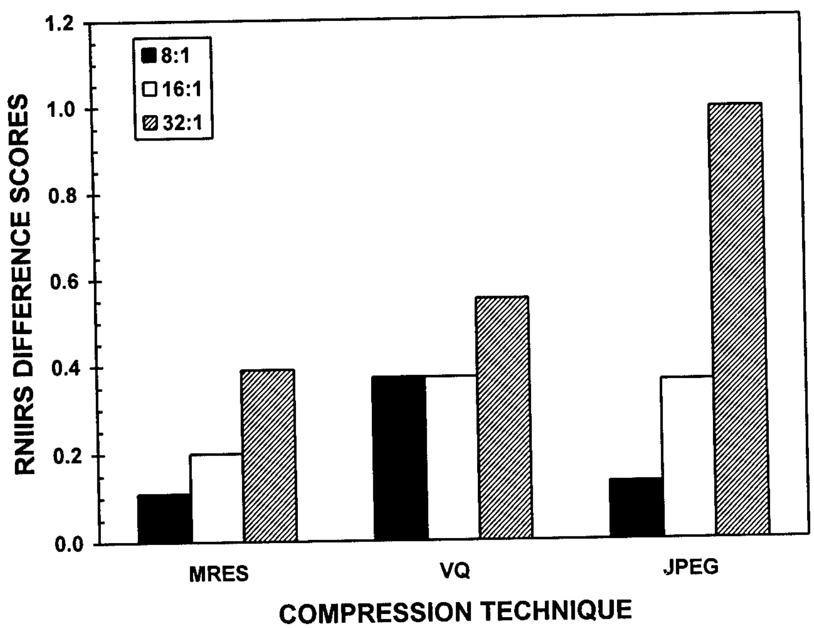

Figure 4. Mean RNIIRS difference scores at ratios of 8:1, 16:1, and 32:1 for each compression technique.

Hits.

Mean difference scores for percentages of hits by resolution, compression technique, and compression ratio are tabulated in Table 6. As with the RNIIRS ratings, a value of 0 indicates that there were no differences in percentages of hits between the uncompressed and compressed imagery. Positive values signify that the hits were greater for the uncompressed imagery, whereas negative values indicate that the hits were greater for the compressed imagery. The means in Table 6 reveal that the hits were consistently greater for the uncompressed imagery, with one exception (JPEG at a ratio of 8:1 for the fine resolution). The difference in hits between uncompressed and compressed imagery was greater at the coarse as opposed to the fine resolution. With respect to compression technique, the difference in hits for uncompressed versus compressed imagery was lowest for MRES and highest for VQ. Finally, the difference score tended to increase as compression ratio increased.

Table 6

Means and Standard Deviations (in Parentheses) of Difference Scores for Percentages of Hits

	COARSE RESOLUTION				FINE RESOLUTION				
	8:1	16:1	32:1	Mean	8:1	16:1	32:1	Mean	MEAN
JPEG	1.2 (0.9)	2.9 (1.2)	7.6 (3.3)	3.9 (1.1)	-0.1 (2.4)	0.4 (1.2)	2.6 (2.1)	1.0 (1.6)	2.3 (0.9)
MRES	0.8 (2.5)	1.6 (2.4)	5.2 (3.5)	2.5 (2.0)	0.2 (3.5)	1.0 (2.2)	3.5 (2.4)	1.6 (1.8)	2.0 (1.3)
VQ	2.5 (1.8)	6.1 (2.1)	7.9 (2.5)	5.5 (1.7)	3.9 (3.0)	3.4 (1.9)	3.4 (1.9)	3.6 (1.6)	4.5 (1.2)
Mean	1.5 (1.3)	3.6 (1.3)	6.9 (2.6)		1.3 (2.7)	1.6 (1.6)	3.1 (1.9)		
Mean	4.0 (1.4)				2.0 (1.6)				2.9 (1.0)

The analysis of variance of the scores in Table 6 revealed that the main effects for compression technique, $\underline{F}(2,10) = 20.48$, $\underline{p} < .0003$, and compression ratio, $\underline{F}(2,10) = 19.04$, $\underline{p} < .0013$, were significant. However, the main effect for resolution did not attain statistical significance, $\underline{F}(1,5) = 3.92$, $\underline{p} > .05$. With respect to the effect for compression technique, post hoc correlated t-tests revealed that JPEG and MRES differed significantly from VQ, but not from each other. Follow-up tests for compression ratio indicated that the ratios of 8:1 and 16:1 differed from 32:1, but not from each other. Of the two-way and three-way interactions, only the interaction between resolution and compression technique was statistically significant, $\underline{F}(2,10) = 6.25$, $\underline{p} < .018$.

The Resolution x Compression Technique interaction is depicted in Figure 5. The main plot in Figure 5 portrays the mean difference scores for percent hits for MRES, VQ, and JPEG at the coarse and fine resolutions. The inset shows the resolution effect for each compression technique (i.e., the difference between the scores at the coarse and fine resolutions for each technique). As can be seen in the figure, the difference scores were consistently larger at the coarse resolution than at the fine resolution for all three compression techniques. However, the difference between the coarse and fine resolutions was smaller for MRES than for either JPEG or

VQ. That is, the effect of resolution on the difference scores for percent hits was lowest for MRES and highest for JPEG and VQ. A post hoc t-test revealed that this effect was statistically significant, $t(10) = 2.53$, $p < .05$.

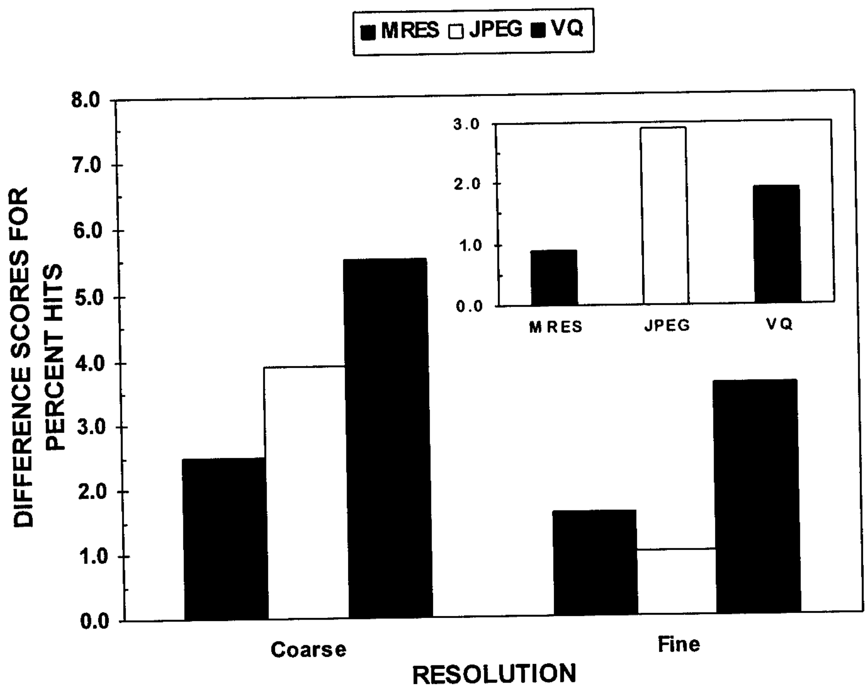

Figure 5. Mean difference scores in percentages of hits for MRES, JPEG, and VQ at the coarse and fine resolutions. (Note: The inset shows the difference between the resolutions for each compression technique.)

False alarms.

Mean difference scores for the number of false alarms by resolution, compression technique, and compression ratio appear in Table 7. As with the other difference scores, a value of 0 indicates no difference between uncompressed and compressed imagery. Positive values indicate that more false alarms were committed with the uncompressed imagery, and negative values signify that more false alarms were made with the compressed imagery. With respect to resolution, the means in Table 7 indicate that more false alarms were made in response to the original imagery as opposed to the compressed imagery at the coarse resolution, but not at the fine resolution. At the fine resolution, slightly more false alarms were made with compressed imagery. For all three compression techniques, there were more false alarms overall with uncompressed imagery than with compressed imagery; however, this effect was primarily the result of the scores at the coarse resolution. Further, the difference scores overall were lowest for MRES and highest for JPEG. This trend signifies that (1) the quantity of false alarms for uncompressed and compressed imagery was most similar for MRES, and (2) the JPEG technique tended to suppress false alarm occurrences more than the other techniques. Inspection of the difference scores by compression ratio reveals no consistent patterns.

Table 7

Means and Standard Deviations (in Parentheses) of Difference Scores for Number of False Alarms

	COARSE RESOLUTION				FINE RESOLUTION				
	8:1	16:1	32:1	Mean	8:1	16:1	32:1	Mean	MEAN
JPEG	0.73 (0.64)	0.60 (0.59)	0.96 (0.63)	0.76 (0.57)	-0.06 (0.08)	-0.03 (0.16)	0.06 (0.08)	-0.01 (0.08)	0.35 (0.25)
MRES	0.57 (0.35)	-0.04 (0.38)	0.32 (0.44)	0.28 (0.21)	-0.11 (0.20)	-0.08 (0.15)	-0.10 (0.12)	-0.09 (0.11)	0.08 (0.06)
VQ	0.40 (0.78)	0.32 (0.36)	0.28 (0.34)	0.34 (0.42)	-0.02 (0.07)	-0.04 (0.24)	0.00 (0.13)	-0.02 (0.14)	0.14 (0.16)
Mean	0.57 (0.53)	0.30 (0.29)	0.52 (0.40)		-0.06 (0.08)	-0.05 (0.18)	-0.01 (0.06)		
Mean				0.46 (0.39)				-0.04 (0.10)	0.19 (0.14)

The outcomes of the analysis of variance indicated that the main effects for resolution, $F(1,5) = 6.92$, $p < .0465$, and compression technique, $F(2,10) = 8.07$, $p < .018$, were significant, as was the Resolution x Compression Technique interaction, $F(2,10) = 9.32$, $p < .01$. None of the remaining main effects or interactions was statistically significant ($p > .05$). With respect to the main effect for compression technique, post hoc correlated t-tests indicated that MRES and VQ differed significantly from JPEG, but not from each other.

The Resolution x Compression Technique interaction is portrayed in Figure 6. The main plot in the figure depicts the mean difference scores for false alarms for MRES, VQ, and JPEG at the coarse and fine resolutions. The inset shows the resolution effect for each technique (i.e., the difference between the scores at the coarse and fine resolutions for each technique). Given that positive difference scores signify more false alarms with original rather than compressed imagery, the plots in Figure 6 reveal that image compression suppressed false alarm responses at the coarse resolution, but not at the fine resolution. However, the magnitude of the difference between the two resolutions was smaller for MRES and VQ than for JPEG. This effect was statistically significant, $t(10) = -3.55$, $p < .05$.

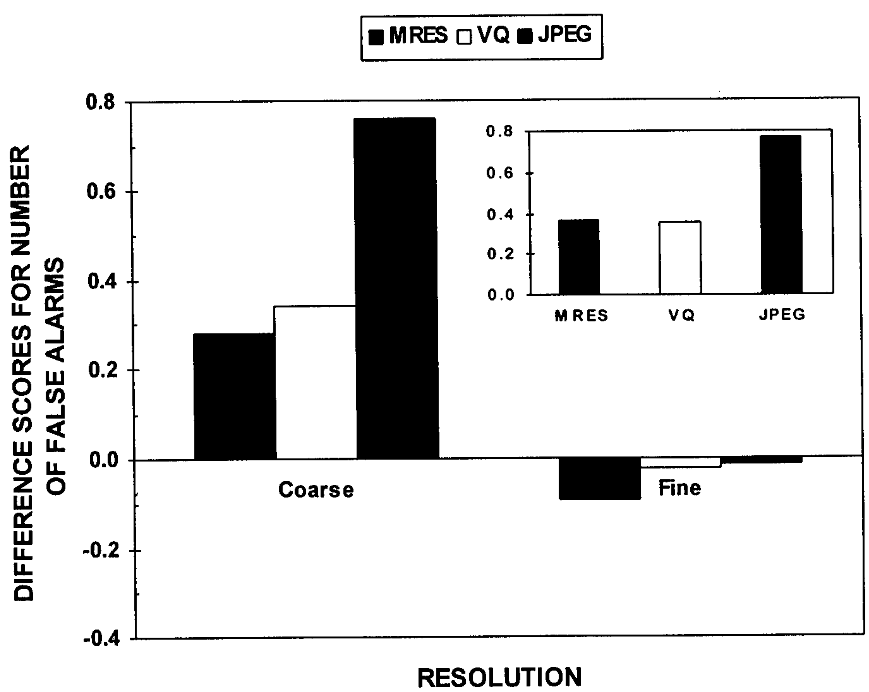

Figure 6. Mean difference scores for the number of false alarms for each compression technique at the coarse and fine resolutions. (Note: The inset shows the difference between the resolutions for each compression technique.)

DISCUSSION

The present study was designed to examine the effects of image compression on the interpretability of SAR imagery by analyzing both subjective and objective measures. Whereas previous investigators have examined either objective or subjective measures of effectiveness, they have often not explored both types of indices simultaneously. Further, many investigations of the effects of compression on subjective measures have been conducted with only a single type of compression technique or with techniques that may not be ideally suited for SAR imagery. In this study, the three different compression techniques (MRES, JPEG, and VQ) that have previously been identified as the most appropriate for SAR imagery were compared in terms of both subjective and objective measures of effectiveness at various compression ratios and image resolutions. The set of original uncompressed imagery consisted of 19 coarse spot scenes and 22 fine spot scenes. Each unique scene was presented once in its original format and nine additional times in versions that had previously been subjected to each of the three compression techniques at each of three compression ratios (8:1, 16:1, and 32:1). Six image analysts rated the interpretability of each image by means of the RNIIRS, a subjective rating scale used to assess the quality of tactical imagery. Subsequently, they identified the location of each target vehicle in the image. The primary measures of interest were the analysts' RNIIRS ratings and their correct and incorrect designation responses (hits and false alarms, respectively). For each dependent variable, difference scores representing the difference in value for comparable uncompressed and compressed images were computed and analyzed.

Subjective Measures of the Effects of Image Compression

Analyses of the difference scores for RNIIRS ratings were used to evaluate the effects of image compression on the analysts' subjective estimates of the interpretability of each image. First, these analyses indicated that the interpretability of coarse resolution imagery was affected by image compression to a greater extent than was fine resolution imagery. Ratings for compressed coarse resolution imagery were significantly lower than those for the original imagery by .56 RNIIRS category units, whereas the difference for the fine resolution imagery was only .24 RNIIRS category units. As described earlier, a difference of .5 on the RNIIRS would be sufficient to reduce either the level of exploitation (e.g., from recognition to detection) or the analyst's decision confidence (e.g., from "probable" to "possible"). Therefore, these outcomes suggest that the compression of coarse resolution imagery can be expected to interfere significantly with the analyst's ability to utilize the imagery. On the other hand, fine resolution imagery of the type used in the current study may be much less susceptible to the degrading effects of image compression. The present results further imply that the interpretability of imagery at even finer resolutions than those employed here may remain relatively unaffected, though this possibility has yet to be tested.

Second, the compression technique that maintained the highest RNIIRS ratings at both the coarse and fine resolutions was MRES. Regardless of which algorithm was applied, the impact of image compression was minimal at the fine resolution (a difference of less than .35 RNIIRS category units for all three techniques). At the coarse resolution, on the other hand, MRES was the only technique that continued to exhibit a relatively small impact. The difference in RNIIRS ratings between the original imagery and MRES was only .35, whereas the differences for JPEG and VQ were greater than .60. These results imply that if one wishes to use image compression without significantly affecting the subsequent interpretability and utility of SAR imagery, MRES will be the most suitable technique.

Third, the degrading effects of image compression on image interpretability increased as the compression ratio increased, though the nature of the effect differed for each technique. In general, RNIIRS difference scores increased as the compression ratio increased for all three techniques, but were smallest for the MRES technique. That is, the subjective ratings for MRES were most comparable to those for the uncompressed imagery, differing by less than .5 RNIIRS category units at even the highest compression ratios. Ratings for VQ were less comparable to those for the original imagery, but were fairly similar at each compression ratio. Finally, for JPEG, there was a pronounced degradation in the RNIIRS rating at the ratio of 32:1, the highest level examined in the current study.

These outcomes parallel the differences that were visually apparent among the three techniques as compression ratio increased. Specifically, MRES images were generally not visually distinguishable from the original imagery, even at the highest compression ratio of 32:1. VQ images tended to appear somewhat grainy or blocky across all compression ratios. Finally, JPEG images appeared similar to the original imagery at ratios of 8:1 and 16:1, but extremely blocky at the highest ratio. There was a noticeable discrepancy between the near-normal visual appearance of JPEG images at ratios of 8:1 and 16:1 and the blockiness at the ratio of 32:1. As described in the introduction, the application of the JPEG algorithm to SAR imagery often produces such visible discontinuities as a consequence of the low spatial correlations in SAR images. The parallels between the statistical outcomes for the image interpretability ratings and the visual appearance of the images indicate that whereas MRES is acceptable at compression ratios as high as 32:1, JPEG is not. The present results further imply that only MRES may perform suitably at even higher ratios than those included here; however, these expectations have not yet been tested empirically.

In summary, examination of the analysts' estimates of image quality in the current study indicated that (1) coarse resolution imagery is more susceptible to the degrading effects of image compression, and (2) MRES is preferable to JPEG and VQ because it produces the least reduction in image interpretability. Given that previous studies involving subjective ratings have assessed only VQ (Penrod & Kuperman, 1993; Vaughn, 1987), these outcomes provide critical data regarding the relative efficacy of various alternative compression algorithms in terms of image interpretability and utility.

Objective Measures of the Effects of Image Compression

In order to arrive at a more conclusive determination of the most effective compression technique for SAR imagery, the subjective RNIIRS ratings were supplemented with analyses of objective measures of performance effectiveness. These objective measures included correct designations of target vehicles (hits) and incorrect designations of nontarget objects as targets (false alarms). Whereas the RNIIRS ratings represented the analysts' opinion regarding the utility of the imagery, the performance measures served as indicators of the actual utility of the imagery for detection purposes. In making their target designations, the image analysts were required only to detect the presence of each target vehicle and not to further classify or identify it in any way.

Hits.

The principal outcome regarding the analysts' correct designations is the finding that MRES supported target acquisition performance to the same extent as the original imagery, regardless of image resolution. Although overall hits, collapsing across resolution, were equally high for both

MRES and JPEG; only MRES had the highest hits at *both* the coarse and fine resolutions, as revealed by inspection of the Compression Technique x Resolution interaction. The hits achieved with the JPEG imagery differed from that with the uncompressed imagery by only one percentage point at the fine resolution; however, this difference increased to nearly four percentage points at the coarse resolution. The comparable scores for MRES, 1.6 and 2.5 percentage points, revealed that it produced the least interference with target detection performance at both resolutions. This trend is consistent with the finding that RNIIRS ratings for the compressed imagery were also most similar to those for the uncompressed imagery at both the coarse and fine resolutions when the imagery had been subjected to the MRES algorithm. Thus, the results pertaining to the hits, an objective indicator of target detection performance, further validate the conclusion derived from inspection of the subjective estimates of image utility; i.e., that MRES is more suitable for SAR imagery than either JPEG or VQ.

<u>False alarms.</u>

Analyses of another objective indicator of target detection performance, the number of false alarm responses, revealed that the general effects of image compression on false alarm occurrences differed depending on the resolution. False alarms were similar in the context of uncompressed and compressed imagery at the fine resolution but were paradoxically greater with the original imagery than with the compressed imagery at the coarse resolution. This outcome may stem from the tendency for image analysts to adopt a more conservative approach with compressed imagery, particularly when the image resolution is relatively coarse to begin with. That is, they may respond only to the most obvious target-like objects in the compressed images and exclude those objects that could potentially be targets but are less definite, an approach that would result in fewer false alarms. On the one hand, the finding that image compression reduces false alarms implies that less time is wasted on the pursuit of nontargets. On the other hand, however, this increased caution may result in the omission of genuine targets that are not readily apparent. Because previous investigations of image compression have not included reports of false alarm occurrences, the reliability of these trends can only be determined through additional research.

Differences in false alarms were also apparent with respect to image compression technique. JPEG suppressed false alarms to the greatest extent, an outcome that may stem from the blocky appearance of images submitted to a JPEG compression. In general, the target vehicles remain visible, but the blockiness obscures other details in the scene, which can exacerbate the task of deciding whether other "target-like" objects are indeed targets. In accordance with these observations, the hits for JPEG were nearly as high as those achieved with the original imagery, but the false alarms were lower. MRES and VQ also suppressed false alarms, but not to the extent that JPEG did. In fact, the difference in false alarms between the uncompressed and compressed imagery approached zero for the MRES and VQ techniques. Further, these two techniques exhibited greater stability in false alarms between the coarse and fine resolutions than did JPEG, as revealed by the Resolution x Compression Technique interaction. Therefore, in terms of false alarm occurrences, MRES and VQ should be regarded as the most suitable techniques since they were most comparable to the original imagery. In summary, the results pertaining to observers' false alarms provide further evidence in support of the conclusions derived from examination of their RNIIRS ratings and correct designation responses. Namely, the three measures of image interpretability and utility examined in the present study consistently indicated that MRES is superior to both JPEG and VQ for the compression of SAR imagery.

Future Research

The outcomes of the current study indicate that the efficacy of each compression technique should be explored further as a function of compression ratio and resolution. Specifically, whereas both subjective and objective measures of effectiveness revealed MRES to be the most suitable technique for SAR imagery, this conclusion is technically applicable only to the ratios and resolutions that were examined in the current investigation. The nature of the results herein suggests that MRES may continue to be more effective than its competitors at compression ratios that exceed 32:1 and that the type of compression technique may become less consequential at finer resolutions. However, these expectations have yet to be verified empirically. Accordingly, future research should focus on the behavior of MRES, JPEG, and VQ at compression ratios greater than 32:1 and at finer resolutions than were examined here.

Future studies should also strive to include images that vary widely in terms of the level of performance they might be expected to support (e.g., detection, classification, identification) in order to foster more complete use of the Radar National Imagery Interpretability Rating Scale. In the current investigation, analysts' ratings tended to cluster at the low end of the scale. Mean

RNIIRS ratings ranged from 1.97 to 3.76, signifying that most if not all of the images were suitable only for target detection. Additional differences among the techniques may emerge with imagery in which classification or identification is possible (i.e., RNIIRS ratings toward the upper end of the scale).

CONCLUSIONS

Of the three techniques, MRES was most similar to uncompressed imagery in terms of both subjective and objective measures of effectiveness. Its image interpretability ratings were highest at both resolutions and even at compression ratios as high as 32:1. The percentage of hits with MRES differed only slightly from those with the original imagery at both the coarse and fine resolutions. Finally, the false alarm occurrences with MRES were also similar to the quantity with the original imagery. These outcomes indicate that MRES is currently a more suitable technique for SAR image compression than either JPEG or VQ.

REFERENCES

ASCC (Air Standardization Coordinating Committee) (10 July 1978). Image Interpretability Rating Scale, AIR STD 101/11.

Crocker, S. C., & Cox, J. L. (1995, June). The Impact of Image Compression on SAR ATC Performance. Paper presented at the 40th ATR Working Group (ATRWG) meeting, Colorado Springs, CO.

Gersho, A., & Gray, R. (1992). Vector quantization and signal compression. Boston: Kluwer Academic Publishers.

Gray, R. M. (1984, April). Vector quantization. IEEE ASSP Magazine, 1, 4-29.

Huynh, H., & Feldt, L. S. (1970). Conditions under which mean square ratios in repeated measures designs have exact F-distributions. Journal of the American Statistical Association, 65, 1582-1589.

Huynh, H., & Feldt, L. S. (1976). Estimation of the Box correction for degrees of freedom from sample data in randomized block and split-plot designs. Journal of Educational Statistics, 1, 69-82.

Itek Optical Systems (1984). IIRS: Image Interpretability Rating Scale, Itek C^3I Systems Bulletin IL-2.

Lan, L. (1995). On the compression behavior of the JPEG baseline image coder. Proceedings of SPIE, The International Society for Optical Engineering: Visual Communications and Image Processing '95, 2501, 573-582.

Lewis, A. (1993). Chicken Little, Grayling 1992. Final Ground Truth Data Report (Report No. 240200-12-X). System Evaluation Program Area.

Mallat, S. G. (1989). A theory for multiresolution signal decomposition: The wavelet representation. IEEE Transactions on Pattern Analysis and Machine Intelligence, 11, 674-693.

NITF (1992). National Imagery Transmission Format (Version 2.0) for the National Imagery Transmission Format Standard, Document ICS-1517/92 Joint Pub-XXX-NITF prepared by NITFS Technical Board, OASD, March 31, 1992.

Pennebaker, W. B., & Mitchell, J. L. (1993). JPEG still image data compression standard. New York: Van Nostrand Reinhold.

Penrod, T. D., & Kuperman, G. G. (1993). Image Quality Analysis of Compressed Synthetic Aperture Radar Imagery (Report No. AL/CF-TR-1993-0156). Wright-Patterson Air Force Base, OH: Armstrong Laboratory.

Rabbani, M., & Jones, P. W. (1991). Digital image compression techniques. Bellingham, WA: SPIE Optical Engineering Press.

Vaughn, V. (1987). A Comparison of Three Image Compression Algorithms Based on Photo Interpretation. Salt Lake City, UT: Unisys.

Wallace, G. (1991). The JPEG still picture compression standard. Commun. ACM, 34, 31-45.

Wharton, J. J., Gorman, J. D., Werness, S. A., & Gleason, J. M. (1993). ARDA Phase I Compression Study: The Impact on ATC Performance.

Wilson, D. L., & Kuperman, G. G. (1989). Effects of ATARS Image Compression on Image Interpretability Ratings (Report No. AAMRL-TR-89-032). Wright-Patterson Air Force Base, OH: Armstrong Aerospace Medical Research Laboratory.

Withman, R. L., Cates, D., & Kotz, R. (1994). Effects of Data Compression on a Morphological Target Screening Algorithm for SAR Imagery. Wright-Patterson Air Force Base, OH: Wright Laboratory.

APPENDIX

RADAR NATIONAL IMAGERY INTERPRETABILITY RATING SCALE

Radar National Imagery Interpretability Rating Scale

Rating Level 0

Interpretability of the imagery is precluded by obscuration. Degradation, or very poor resolution.

Rating Level 1

Detect the presence of aircraft dispersal parking areas.
Detect a large cleared swath in a densely wooded area.
Detect, based on presence of piers and warehouses, a port facility.
Detect lines of transportation (either road or rail, but do not distinguish between).

Rating Level 2

Detect the presence of large (e.g., BLACKJACK, CAMBER, COCK: 707, 747) bombers or transports.
Identify large phased array radars (e.g., HEN HOUSE, DOG HOUSE) by type.
Detect a military installation building pattern and site configuration.
Detect road pattern, fence and hardstand configuration at SSM launch sites (missile silos, launch control silos), within a known ICBM complex.
Detect large non-combatant ships (e.g., freighters or tankers) at a known port facility.
Identify athletic stadiums.

Rating Level 3

Detect medium sized aircraft.
Identify an ORBITA site on the basis of a 12 meter dish antenna normally mounted on a circular building.
Detect vehicle revetments at a ground forces facility.
Detect vehicles/pieces of equipment at a SAM, SSM, or ABM fixed missile site.
Determine the location of the superstructure (e.g., fore, amidships, aft) on a medium-size freighter.
Identify a medium size (approx. six track) railroad classification yard.

Rating Level 4

Distinguish between large rotary-wing and medium fixed-wing aircraft (e.g., HALO helicopter vs. CRUSTY transport).
Detect recent cable scars between facilities or command posts.
Detect individual vehicles in a row at a known motor pool.
Distinguish between open and closed sliding roof areas on a single bay garage at a mobile missile base.
Identify square bow shape of ROPUCHA class (LST).
Detect all rail/road bridges.

Rating Level 5

Count all medium helicopters (e.g., HIND, HIP, HAZE, HOUND, PUMA, WASP).
Detect deployed TWIN EAR antenna.
Distinguish between river crossing equipment and medium/heavy armored vehicles by size and shape (e.g., MTU-20 vs. T-62 MBT).
Detect missile support equipment at an SS-25 RTP (e.g., TEL, MSV).
Distinguish bow shape and length/width differences of SSNs.
Detect the break between railcars (count railcars).

Rating Level 6

Distinguish between variable and fixed-wing fighter aircraft (e.g., FENCER vs. FLANKER).
Distinguish between the BAR LOCK and SIDE NET antennas at a BAR LOCK/SIDE NET acquisition radar site.
Distinguish between small support vehicles (e.g., UAZ-69, UAZ-169) and tanks (e.g., T-72, T80).
Identify SS-24 launch triplet on a known location.
Distinguish between the raised helicopter deck on a KRESTA II (CG) and the helicopter deck with main deck on a KRESTA I (CG).
Identify a vessel by class when singly deployed (e.g., YANKEE I, DELTA I, KRIVAK II FFG).
Detect cargo on a railroad flatcar or in a gondola.

Rating Level 7

Identify small fighter aircraft by type (e.g., FISHBED, FITTER, FLOGGER).
Distinguish between electronics van trailers (without tractor) and van trucks in garrison.
Distinguish by size and configuration between a turreted, tracked APC and a medium tank (e.g., BMP-1/2 vs. T-64).
Detect a missile on the launcher in a SA-2 launch revetment.
Distinguish between bow mounted missile system on KRIVAK I/II and bow mounted gun turret on KRIVAK III.
Detect road/street lamps in an urban, residential area or military complex.

Rating Level 8

Distinguish the fuselage difference between a HIND and a HIP helicopter.
Distinguish between the FAN SONG E missile control radar and the FAN SONG F based on the number of parabolic dish antennas (three vs. one).
Identify the SA-6 transloader when other SA-6 equipment is present.
Distinguish limber hole shape and configuration differences between DELTA I and YANKEE I (ISBNs).
Identify the dome/vent pattern on rail tank cars.

Rating Level 9

Detect major modifications to large aircraft (e.g., fairings, pods, winglets).
Identify the shape of antennas on EW/GCI/ACQ radars as parabolic, parabolic with clipped corners, or rectangular.
Identify, based on presence or absence of turret, size of gun tube, and chassis configuration wheeled or tracked PACs by type (e.g., BTR-80, BMP-1/2, MT-LB, M113).
Identify the forward fins on an SA-3 missile.
Identify individual hatch covers of vertically launched SA-N-6 surface-to-air system.
Identify trucks as a cab-over-engine or engine-in-front.

GLOSSARY

APC	Armored Personnel Carrier
ASCC	Air Standardization Coordinating Committee
ATC/ATR	Automatic Target Cuer/Automatic Target Recognizer
BDA	Bomb Damage Assessment
Bit	Binary Digit
CIWAL	Crew-Aiding and Information Warfare Analysis Laboratory
DCT	Discrete Cosine Transform
DoD	Department of Defense
DPCM	Differential Pulse Code Modulation
GOB	Ground Order Battle
GRD	Ground Resolved Distance
IIRS	Image Interpretability Rating Scale
ISO	International Standardization Organization
JPEG	Joint Photographic Experts Group
M	Mean
MRES	Multiresolution Encoding
MSE	Mean Square Error
NAIC	National Air Intelligence Center
NITFS	National Imagery Transmission Format Standards
PTSVQ	Pruned Tree-Structured Vector Quantization
RNIIRS	Radar National Imagery Interpretability Rating Scale
SAR	Synthetic Aperture Radar
SD	Standard Deviation
SNR	Signal-to-Noise Ratio
VQ	Vector Quantization